
Subject: Re: [RFC] [PATCH 2/2] namespace enter: introduce sys_hijack (v3)

Posted by [serue](#) on Mon, 10 Sep 2007 19:17:43 GMT

[View Forum Message](#) <> [Reply to Message](#)

Quoting Paul Menage (menage@google.com):

> On 9/4/07, Serge E. Hallyn <serue@us.ibm.com> wrote:

> > We could of course have the ns_container subsystem do that. The

> > ns_container generally stick around until the admin does a manual rm on

> > its directory, so this way we could keep the nsproxy around.

>

> So how about taking sys_hijack() even further and integrate it with

> control groups too? So when you do sys_hijack() (or maybe an

> alternative name would be sys_fork_in()?) you create a task that

> inherits all the control groups of the target task, as well as the

> namespaces.

>

> Paul

Sorry don't know why i haven't replied to this.

Good point. I see container_fork(p) takes the container from current.

I can change that to container_fork(src, dest) in my next posting.

Is there any reason why we wouldn't want to do that? For instance a container admin could impose some restrictions which would keep the host admin from doing something through sys_hijack()? (Not sure that's a big worry since the restrictions would apply to the container admin as well)

thanks,

-serge

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
